

Claims

- [c1] 1.A computer system comprising:
a host entity for issuing an IO request;
a redundant external storage virtualization controller (SVC) pair coupled to the host entity for performing an IO operation in response to the IO request issued by the host entity comprising a first and a second external storage virtualization controller coupled to the host entity;
and
a plurality of physical storage devices for providing storage to the computer system, each of the physical storage devices coupled to the redundant storage virtualization controller pair through a point-to-point serial signal interconnect;
wherein when the second storage virtualization controller is not on line, the first storage virtualization controller will automatically take over the functionality originally performed by the second storage virtualization controller.
- [c2] 2.The computer system of claim 1, wherein for each of the physical storage devices, the computer system further comprises an access control switch coupled between

the physical storage device and the redundant storage virtualization controller pair for selectively switching the connection of the physical storage device to the redundant SVC pair between the first and the second storage virtualization controller.

- [c3] 3.The computer system of claim 1 wherein in the redundant storage virtualization controller pair, each of the storage virtualization controllers further comprises:
a host-side IO device interconnect controller coupled to the host entity;
a central processing circuitry coupled to the host-side IO device interconnect controller for performing the IO operation in response to the IO request issued by the host entity; and
a plurality of device-side IO device interconnect controllers coupled to the central processing circuitry;
wherein each of the physical storage devices is coupled to the device-side IO device interconnect controllers through the point-to-point serial signal interconnect.
- [c4] 4.The redundant storage virtualization computer system of claim 1 wherein the point-to-point serial signal interconnect is a Serial ATA IO device interconnect.
- [c5] 5.A redundant storage virtualization subsystem for providing storage to a host entity, comprising:

a redundant external storage virtualization controller (SVC) pair for coupling to the host entity for performing an IO operation in response to an IO request issued by the host entity comprising a first and a second storage virtualization controller for coupling to the host entity; and

a plurality of physical storage devices (PSDs) for providing storage to the host, each of the physical storage devices coupled to the redundant storage virtualization controller pair through a point-to-point serial signal interconnect;

wherein when the second storage virtualization controller is not on line, the first storage virtualization controller will automatically take over the functionality originally performed by the second storage virtualization controller.

[c6] 6.The redundant storage virtualization subsystem of claim 5 further comprising an access control switch coupled between a said the physical storage device and the redundant storage virtualization controller pair for selectively switching the connection of the physical storage device to the redundant storage virtualization controller pair between the first and the second storage virtualization controller.

- [c7] 7. The redundant storage virtualization subsystem of claim 5, wherein a said PSD is received in a canister removably attached to the redundant storage virtualization subsystem.
- [c8] 8. The redundant storage virtualization subsystem of claim 7, wherein said PSD is a SATA PSD.
- [c9] 9. The redundant storage virtualization subsystem of claim 7, wherein said PSD is a PATA PSD.
- [c10] 10. The redundant storage virtualization subsystem of claim 5 further comprising an access control switch coupled between a said physical storage device and the redundant storage virtualization controller pair for selectively allowing patching through of the serial signal of the physical storage device to and from the first SVC when in a first patching state of said access control switch and to and from the second SVC when in a second patching state of said access control switch.
- [c11] 11. The redundant storage virtualization subsystem of claim 10, wherein an access ownership arbitration mechanism is provided between said SVC pair and said access control switch to control the patching state of said access control switch.
- [c12] 12. The redundant storage virtualization subsystem of

claim 11, wherein said access ownership arbitration mechanism comprises a pair of access request signal lines coupled between said SVC pair for each complementary pair of device-side IO device interconnect from said SVC pair to said access control switch; said first SVC being active on a first of said access request signal line pair and passive on a second of said access request signal line pair; said second SVC being active on said second and passive on said first of said access request signal line pair; and said SVC pair each being capable of issuing access request signal on its own said active access request signal line, and reading a requesting state of said access request signal and identifying a change of said requesting state since previous reading on its own said passive access request signal line.

[c13] 13. The redundant storage virtualization subsystem of claim 11, wherein said access ownership arbitration mechanism comprises an access ownership detecting mechanism to determine if access ownership is possessed by a said SVC.

[c14] 14. The redundant storage virtualization subsystem of claim 11, wherein said access ownership arbitration mechanism comprises an access ownership granting mechanism to grant access ownership when said access ownership is requested by a said SVC.

[c15] 15. The redundant storage virtualization subsystem of claim 11, wherein said access ownership arbitration mechanism comprises a first and a second access ownership arbitration circuit (AOAC) each coupled to said first and second SVCs and said access control switch, and wherein if said first SVC issues a first access ownership request signal received by said first AOAC, access ownership will be granted to said first SVC when said second SVC does not possess the access ownership, and if said second SVC issues a second access ownership request signal received by said second AOAC, access ownership will be granted to said second SVC when said first SVC does not possess the access ownership.

[c16] 16. The redundant storage virtualization subsystem of claim 15, further comprises an access ownership determining mechanism whereby when said first and said second SVC concurrently issues said first and second access ownership request signal to said first and said second AOAC, respectively, access ownership will be granted to a predetermined one of said SVC pair.

[c17] 17. The redundant storage virtualization subsystem of claim 5 wherein in the redundant storage virtualization controller pair, each of the storage virtualization controllers further comprises:

at least one host-side IO device interconnect controller for coupling to the host entity;
a central processing circuitry coupled to the host-side IO device interconnect controller for performing the IO operation in response to the IO request issued by the host entity; and
a plurality of device-side IO device interconnect controller coupled to the central processing circuitry;
wherein each of the physical storage devices is coupled to the device-side IO device interconnect controllers through the point-to-point serial signal interconnect.

[c18] 18. The redundant storage virtualization subsystem of claim 17, wherein said at least one host-side IO device interconnect controller each comprises at least one host-side port, a said host-side port of the first SVC and a said host-side port of said second SVC constitute a complementary port pair coupled to a same said host-side IO device interconnect.

[c19] 19. The redundant storage virtualization subsystem of claim 17, wherein said at least one host-side IO device interconnect controller each comprises at least one host-side port, and a said host-side port of the first SVC and a said host-side port of said second SVC constitute a complementary port pair each coupled to a different said host-side IO device interconnect.

- [c20] 20. The redundant storage virtualization subsystem of claim 17, wherein a logical media unit is redundantly presented to the host entity by said redundant SVC pair on said complementary port pair.
- [c21] 21. The redundant storage virtualization subsystem of claim 17, wherein a said SVC includes a plurality of host-side IO device interconnect ports in said at least one host-side IO device interconnect controller.
- [c22] 22. The redundant storage virtualization subsystem of claim 21, wherein a logical media unit is redundantly presented to the host entity on more than one of said host-side IO device interconnect port.
- [c23] 23. The redundant storage virtualization subsystem of claim 5, wherein said host-side IO device interconnect controller is of a type selecting from one of the followings: a Fibre Channel controller supporting Fabric, point-to-point, public loop and/or private loop connectivity in target mode, a parallel SCSI controller operating in target mode, ethernet controller supporting the iSCSI protocol operating in target mode, and a serial SCSI controller operating in target mode.
- [c24] 24. The redundant storage virtualization subsystem of claim 5 wherein the point-to-point serial signal inter-

connect is a Serial ATA IO device interconnect.

- [c25] 25. The redundant storage virtualization subsystem of claim 5, wherein when the first storage virtualization controller is not on line, the second storage virtualization controller will automatically take over the functionality originally performed by the first storage virtualization controller.
- [c26] 26. The redundant storage virtualization subsystem of claim 25, further comprising an access control switch for a pair of device-side IO device interconnects connecting to different SVCs and configured in redundant pair, said access control switch being coupled between the redundant storage virtualization controller pair for selectively switching the connection of the physical storage device to one of said redundant storage virtualization controllers; a cooperating mechanism for the SVC pair to cooperatively control a patching state of said access control switch; a monitoring mechanism for each SVC of said SVC pair to monitor status of the other SVC of said SVC pair; and, a state control mechanism for each SVC of said SVC pair to forcibly taking complete control of the other SVC of said SVC pair.
- [c27] 27. The redundant storage virtualization subsystem of claim 5, wherein an inter-controller communication

channel is provided between said redundant SVC pair for communicating state synchronization information.

- [c28] 28. The redundant storage virtualization subsystem of claim 27, wherein a said inter-controller communication channel is an existing IO device interconnect, whereby inter-controller communication exchange is multiplexed with IO requests and associated data.
- [c29] 29. The redundant storage virtualization subsystem of claim 27, wherein a said inter-controller communication channel is a dedicated channel the primary function thereof is to exchange said state synchronization information.
- [c30] 30. The redundant storage virtualization subsystem of claim 27, wherein a said inter-controller communication channel is selecting from one of the following: Fibre, SATA, Parallel SCSI, Ethernet, Serial SCSI, I2C.
- [c31] 31. The redundant storage virtualization subsystem of claim 5, wherein said redundant SVC pair can perform IO request rerouting function.
- [c32] 32. The redundant storage virtualization subsystem of claim 5, wherein said redundant SVC pair can perform PSD access ownership transfer function.

[c33] 33. The redundant storage virtualization subsystem of claim 5, wherein each SVC of said redundant SVC pair includes at least one expansion port for coupling to a second plurality of PSDs through a multiple-device device-side IO device interconnects and each said expansion port on said first SVC has a complementary expansion port on the second SVC.

[c34] 34. The redundant storage virtualization subsystem of claim 33, wherein said second plurality of PSDs each have complementary IO ports in dual-port pair, each connecting to a different one of said complementary expansion port pair.

[c35] 35. The redundant storage virtualization subsystem of claim 33, wherein said complementary expansion port pair are connected to the said second plurality of PSDs through a switch circuit.

[c36] 36. The redundant storage virtualization subsystem of claim 33, wherein each of said complementary expansion port pair has a redundant complement expansion port on the same SVC, and said complementary expansion port and redundant complement expansion port are connected to a different one of complementary IO ports in dual-port pair of a said second plurality of PSDs.

- [c37] 37. The redundant storage virtualization subsystem of claim 36, wherein said complementary expansion port pair are connected to the second plurality of PSDs through a switch circuit.
- [c38] 38. The redundant storage virtualization subsystem of claim 33, wherein a said device-side expansion port and interconnect is of a type selecting from one of the following: Fibre, Parallel SCSI, Expanded Serial ATA, Ethernet, and Serial SCSI.
- [c39] 39. The redundant storage virtualization subsystem of claim 5, wherein said first and second SVC each includes a first and a second expansion port for coupling to a second plurality of PSDs through a multiple-device device-side IO device interconnects and said second plurality of PSDs each have a first and a second IO port forming a dual-ported port pair of said second plurality of PSDs, said first and second expansion ports of said first SVC formed a first redundant complement, said first and second expansion ports of said second SVC formed a second redundant complement, said first expansion ports of said SVCs formed a third redundant complement, said second expansion ports of said SVCs formed a fourth redundant complement, and wherein an interconnect signal line switching mechanism is provided for each said redundant complement between correspond-

ing IO device interconnects to switch connection to said dual-ported port pair.

- [c40] 40. The redundant storage virtualization subsystem of claim 39, wherein said interconnect signal line switching mechanism supports one of the following arrangements:
- (1) direct connecting of a said expansion port on the first SVC to a first PSD IO port of said dual-ported port pair and direct connecting of a said expansion port on the second SVC to a second PSD port of said dual-ported port pair;
 - (2) interconnecting of two expansion ports in said redundant complement to an IO device interconnect in the redundant complement that connects to said first PSD IO port;
 - (3) interconnecting of two expansion ports in said redundant complement to an IO device interconnect in the redundant complement that connects to said second PSD IO port;
 - (4) directly connecting of a said expansion port on the first SVC to said IO device interconnect in the redundant complement that connects to said first PSD IO port;
 - (5) directly connecting of a said expansion port on the first SVC to said IO device interconnect in the redundant complement that connects to said second PSD IO port;
 - (6) directly connecting of a said expansion port on the

second SVC to said IO device interconnect in the redundant complement that connects to said first PSD IO port; and,

(7) directly connecting of a said expansion port on the second SVC to said IO device interconnect in the redundant complement that connects to said second PSD IO port.

- [c41] 41. The redundant storage virtualization subsystem of claim 39, wherein said interconnect signal line switching mechanism supports all of the following arrangements:
- (1) direct connecting of a said expansion port on the first SVC to a first PSD IO port of said dual-ported port pair and direct connecting of a said expansion port on the second SVC to a second PSD port of said dual-ported port pair;
 - (2) interconnecting of two expansion ports in said redundant complement to an IO device interconnect in the redundant complement that connects to said first PSD IO port;
 - (3) interconnecting of two expansion ports in said redundant complement to an IO device interconnect in the redundant complement that connects to said second PSD IO port;
 - (4) directly connecting of a said expansion port on the first SVC to said IO device interconnect in the redundant

complement that connects to said first PSD IO port;
(5) directly connecting of a said expansion port on the first SVC to said IO device interconnect in the redundant complement that connects to said second PSD IO port;
(6) directly connecting of a said expansion port on the second SVC to said IO device interconnect in the redundant complement that connects to said first PSD IO port;
and,
(7) directly connecting of a said expansion port on the second SVC to said IO device interconnect in the redundant complement that connects to said second PSD IO port.

[c42] 42. The redundant storage virtualization subsystem of claim 5, wherein said first SVC of said redundant SVC pair includes a state-defining circuit for forcefully defining externally connected signal lines of said second SVC to a predetermined state.

[c43] 43. The redundant storage virtualization subsystem of claim 5, wherein each SVC of said redundant SVC pair includes a self-killing circuit for forcefully defining externally connected signal lines thereof to a predetermined state.

[c44] 44. An external storage virtualization controller for use in a redundant storage virtualization controller pair, com-

prising:

a host-side IO device interconnect controller for coupling to a host entity;

a central processing circuitry coupled to the host-side IO device interconnect controller for performing an IO operation in response to an IO request issued by the host entity;

a memory coupled to the central processing circuitry;
and

at least one device-side IO device interconnect controller coupled to the central processing circuitry, for performing point-to-point serial signal transmission with a plurality of physical storage devices;

wherein when a second external storage virtualization controller in the redundant storage virtualization controller pair is not on line, said external storage virtualization controller will automatically take over the functionality originally performed by the second external storage virtualization controller.

[c45] 45. The storage virtualization controller of claim 44 wherein the device-side IO device interconnect controller is a Serial ATA IO device interconnect controller comprising a plurality of Serial ATA ports, each for connecting to a said physical storage device through a Serial ATA IO device interconnect.

- [c46] 46. The storage virtualization controller of claim 44 further comprising an off-line detecting mechanism for detecting an off-line state of said second storage virtualization controller.
- [c47] 47. The storage virtualization controller of claim 44 wherein said functionality includes presenting and making available to the host entity accessible resources that were originally presented and made available by said second storage virtualization controller as well as accessible resources that were presented and made available by said storage virtualization controller itself.
- [c48] 48. A method for performing storage virtualization in a computer system having a first and a second external storage virtualization controller, the method comprising: performing an IO operation by the second storage virtualization controller in response to an IO request issued by a host entity of the computer system to access at least one of a plurality of physical storage devices of the computer system in point-to-point serial signal transmission; and
when the second storage virtualization controller is not on line, performing the IO operation by the first storage virtualization controller in response to the IO request issued by the host entity to access said at least one of the

physical storage devices of the computer system in point-to-point serial signal transmission.

[c49] 49. The method of claim 48 wherein the point-to-point serial signal transmission is performed in a format complying with Serial ATA protocol.

[c50] 50. The method of claim 48 wherein said first storage virtualization controller will automatically take over the functionality originally performed by said second storage virtualization controller when the second storage virtualization controller is not on line.

[c51] 51. The method of claim 50 wherein said functionality includes presenting and making available to the host entity accessible resources that were originally presented and made available by said second storage virtualization controller as well as accessible resources that were presented and made available by said first storage virtualization controller itself.

[c52] 52. The method of claim 48, further comprising providing a rerouting mechanism for said SVC pair to perform IO request rerouting function.

[c53] 53. The method of claim 52, wherein said IO request rerouting function is performed by the steps of:
a request initiator of said SVC pair transferring IO re-

quest to an access owner of said SVC pair;
said access owner performing said IO request transferred from said request initiator; whereby data associated with said IO request transferred between said access owner and said PSDs are forwarded over to said request initiator.

[c54] 54. The method of claim 48, further comprising the steps of:
performing an IO operation by the first storage virtualization controller in response to a second IO request issued by said host entity to access said at least one of a plurality of physical storage devices of the computer system in point-to-point serial signal transmission; and
when the first storage virtualization controller is not on line, performing the IO operation by the second storage virtualization controller in response to the IO request issued by the host entity to access said at least one of the physical storage devices in point-to-point serial signal transmission.

[c55] 55. The method of claim 54, further comprising the steps of:
providing an access control switch coupled between a said the physical storage device and the redundant storage virtualization controller pair for selectively allowing patching through of the serial signal of the physical

storage device to and from the first SVC when in a first patching state of said access control switch and to and from the second SVC when in a second patching state of said access control switch;

providing a first signal line being of said first SVC active and of said second SVC passive for issuing a first access request signal from said first SVC to said second SVC;

providing a second signal line being of said second SVC active and of said first SVC passive for issuing a second access request signal from said second SVC to said first SVC;

an access requester of said SVC pair asserting its active signal line to request access ownership to a said PSD from an access owner of said SVC pair;

said access owner deasserting its active signal line to relinquish said access ownership; and

said access requester asserting its active signal line and changing said patching state of said access control switch to acquire said access ownership.

- [c56] 56. The method of claim 55, further comprising the steps of said access owner holding up and queuing up new IO requests for later execution and completing all pending IOs, after the step of said access requester asserting its active signal line and before the step of said access owner deasserting its active signal line.

[c57] 57. The method of claim 54, further comprising the steps of:

providing an access control switch coupled between a said PSD and the redundant storage virtualization controller pair for selectively switching the connection of the PSD to the redundant storage virtualization controller pair between the first and the second storage virtualization controller;

providing a first access ownership arbitration circuit for acquiring access ownership of said access control switch for said first SVC having a first access ownership request (AOR) signal line as a first input line thereof couple to said first SVC, and a first access control switch control (ACSCS) line coupled to said access control switch as a first output line;

providing a second access ownership arbitration circuit for acquiring access ownership of said access control switch for said second SVC having a second access ownership request signal line as a first input line thereof couple to said second SVC, and a second access control switch control signal line coupled to said access control switch as a second output line; whereby

when one SVC of said SVC pair asserts its AOR signal line while the other SVC of said SVC pair having been asserting its ACSCS line, said ACSCS line of said one SVC will

not be asserted until said the other SVC deasserts its said ACSCS line.

- [c58] 58. The method of claim 57, further comprising the steps of:
- providing said first access ownership arbitration circuit a first alternate SVC access ownership request (ASAOR) signal line as a second input line thereof couple to said second SVC;
 - providing said second access ownership arbitration circuit a second alternate SVC access ownership request (ASAOR) signal line as a second input line thereof couple to said first SVC; whereby
 - when a said ASAOR signal line of one SVC of said SVC pair is asserted, said ASAOR signal line of the other SVC of said SVC pair will be asserted unless said first and second SVCs are asserting said AOR signal lines concurrently.

- [c59] 59. The method of claim 58, further comprising the steps of:
- providing an access ownership determining mechanism for granting access ownership of said access control switch to one SVC of said SVC pair when said first and second SVCs are asserting said AOR signal lines concurrently.

- [c60] 60. The method of claim 59, wherein said access ownership determining mechanism grants access ownership of said access control switch to said one SVC of said SVC pair by asserting said ASAOR signal line coupled to said access ownership arbitration circuit for the other SVC of said SVC pair.
- [c61] 61. The method of claim 57, wherein information exchanges associated with access ownership transfer between said SVC pair is communicated as a part of inter-controller communications.
- [c62] 62. The method of claim 48, further comprising providing an access ownership transferring mechanism for one of said SVCs that possessed access ownership of said PSD to transfer said access ownership to the other of said SVCs.
- [c63] 63. The method of claim 62, wherein said access ownership transferring mechanism performing the steps of:
- (a) an access requester of said SVC pair issuing an access request signal to an access owner of said SVC pair for requesting access ownership to a said PSD;
 - (b) said access owner relinquishing said access ownership such that it is not an access owner now; and
 - (c) said access requester acquiring said access ownership and becoming an new access owner of said PSD.

- [c64] 64. The method of claim 63, further comprising the steps of said access owner holding up and queuing up new IO requests for later execution and completing all pending IOs, after the step of (a) said access requester issuing said access request signal to said access owner and before the step of (b) said access owner relinquishing said access ownership.
- [c65] 65. The method of claim 63, wherein in step of (b), said access owner relinquishing said access ownership by modifying a state of an access control switch coupled between said SVC pair and said PSD.
- [c66] 66. The method of claim 63, wherein in step of (c), said access requester acquiring said access ownership by modifying a state of an access control switch coupled between said SVC pair and said PSD.
- [c67] 67. A computer-readable storage medium having a computer program code stored therein that is capable of causing a computer system having a host entity, a first and a second external storage virtualization controller coupled to the host entity and a plurality of physical storage devices coupled to the first and the second storage virtualization controller to perform the steps of: performing an IO operation by the second storage virtu-

alization controller in response to an IO request issued by the host entity to access at least one of the physical storage devices in point-to-point serial signal transmission; and
said first storage virtualization controller automatically performing the IO operation that was originally performed by the second storage virtualization controller in response to the IO request issued by the host entity to access at least one of the physical storage devices in point-to-point serial signal transmission when the second storage virtualization controller is not on line.

[c68] 68.The computer readable medium of claim 67 wherein each of the physical storage devices is coupled to the first and the second storage virtualization controller through a Serial ATA IO device interconnect.